EXHIBIT 4

1. Standards and expectations for use of western blotting to quantify protein content in biological samples

Paul S. Brookes, PhD, Professor of Anesthesiology, University of Rochester Medical Center

Contents:

1. Outline

In this document I will describe the process of western blotting, and accepted standards in the biomedical sciences for the preparation, presentation, and quantitation of western blot data. This is based on my experience with this method for over 20 years, and appropriate citations. After describing the experimental procedure and the necessary quality controls to obtain reliable data, discussion will cover common mistakes, manipulations or other issues that can undermine confidence in conclusions arising from blot data.

2. Purpose of western blotting and method overview

In biomedical science it is often necessary to quantify the amount of a protein and how this responds to a perturbation (e.g., a drug treatment). Western blotting is one method to achieve this, and comprises 3 steps: (i) Proteins in a sample are separated by their molecular weight using gel electrophoresis (SDS-PAGE) and the separated proteins are then transferred to a membrane. (ii) The membrane is probed using an antibody that binds to the protein of interest. (iii) The amount of antibody binding is quantified via chemiluminescence. At each step, quality controls are necessary to ensure the eventual readout reliably reports the amount of protein in each sample. Such controls include use of a *molecular weight ladder* (a mixture of proteins of defined size) to calibrate the gel, and the use of *loading controls* to ensure the same amount of total protein from each sample is loaded onto the gel. A key principle throughout, is that all samples and steps must be treated equally.

(i) Step 1, SDS-PAGE and membrane transfer

Biological samples (e.g., cells, tissues, homogenates) are extracted in liquid form, by boiling them in a solution containing the detergent sodium dodecyl sulfate (SDS) and the chemical β-mercaptoethanol. This causes proteins in the sample to denature (unfold). SDS coats the proteins and imparts a negative charge on them. A protein assay can be performed to quantify the total amount of protein in each sample. However, this is not always possible (as SDS can interfere with some commonly used protein assay methods), so a more common procedure is to use a *loading control* at a later stage. The proteins dissolved in SDS are separated by gel electrophoresis. Two glass plates are separated by a thin spacer (typically 1.5mm thick), and a mixture of acrylamide/bis-acrylamide is poured into the gap, along with agents that cause the acrylamide to polymerize and solidify. An extra layer of gel is then poured and set, with a plastic comb used to create slots in the top of the gel. When the comb is removed, the spaces left behind are called *wells*. The typical gel has 10 or 15 wells, so can be used to analyze 10 or 15 biological samples. The image on the left shows a glass plate assembly with a 10-well comb.

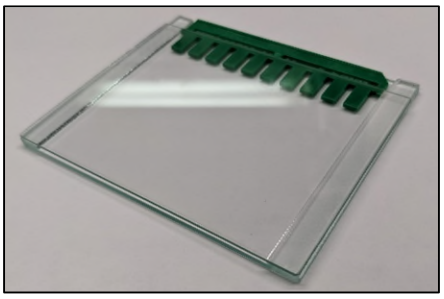

The polymerized gel is placed in a *gel-box* apparatus, comprising a tank filled with a solution containing SDS. The solution conducts electricity, and located within the tank are platinum electrodes. The cathode (negative) is located at the top of the gel, near the wells. The anode (positive) is located near the bottom of the gel. The protein samples are loaded by pipetting into the wells, and the gel box is then connected to a power supply (10-100 Volts DC) for 1-3 hours. The electric charge across the gel causes the proteins (negatively charged with SDS) to migrate down the gel toward the anode. The gel matrix contains pores that act as a sieve to slow down the proteins, such that small proteins move quickly through the gel, while large proteins move slowly. When preparing a gel, the percentage acrylamide determines the pore size… 12-15% gels are used to separate large proteins, 5-10% for smaller proteins. At the end of the run time, the result is a separation of proteins below each well, with large proteins at the top near their starting point, and smaller proteins at the bottom having migrated further. The vertical space below each well is termed a *lane*, and within each lane the proteins appear as horizontal *bands* (see image at left).

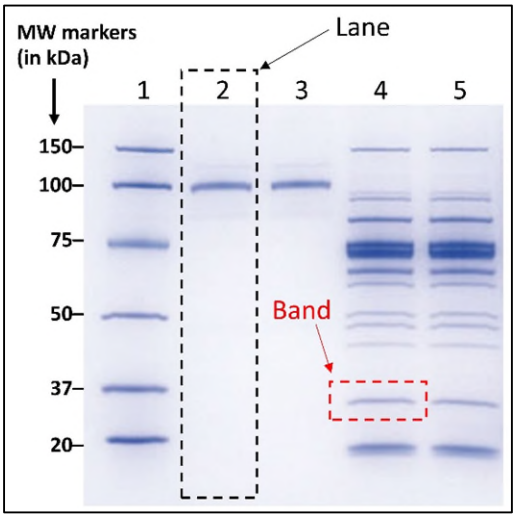

It is usually necessary to calibrate the gel by loading *molecular weight markers* in one of the wells. This is a mixture of proteins of known molecular weight (MW), often tagged with colored dyes so their position on the gel can easily be seen. In the image above, the markers are loaded in lane 1, and their molecular weight (in kiloDaltons, kDa) is annotated alongside the gel. The vertical position of a protein relative to the MW ladder is used to calculate its size (here the red highlighted red band has a mass between 20 and 37 kDa). When probing multiple proteins on a single western blot, it is critical to use a MW ladder so the membrane can be cut into horizontal slices (representing different MW ranges) to probe with different antibodies.

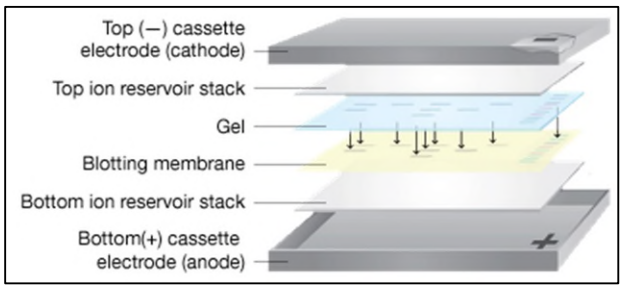

Gels are very fragile, so the next step is to transfer the separated proteins to more stable substrate. The substrate of choice is a membrane of nitrocellulose or PVDF – both materials that bind proteins well. The transfer process (see image at left) is termed *western blotting*, with the resulting *blot* being an imprint of the gel. It is common to mark the membrane (e.g., a nick in one corner) to orient it relative to the original gel at a later stage. An important feature arising from the above process, is that no two gels or membrane blots are identical.

(ii) Step 2, probing the membrane with antibodies

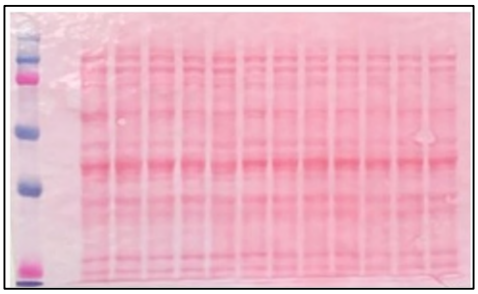

At this stage, the MW marker proteins usually appear on the membrane as colored bands visible to the naked eye, whereas the rest of the membrane appears blank white. To check that the transfer occurred correctly, it is common to stain the membrane with the dye Ponceau red (see image at left), which can then be washed away before processing the blot further. Because the membrane binds protein, and antibodies themselves are proteins, before the antibody incubation step it is necessary to *block* the unoccupied protein binding sites on the membrane, by incubating it with a uniform protein solution. Typically, non-fat dry milk (casein) or bovine serum albumin are used as cheap sources of blocking protein.

The *blocked* membrane is then probed with an antibody that specifically binds to the protein of interest - this is termed the *primary antibody*. Typical antibodies for biomedical science are raised in mouse or rabbit. Mouse antibodies are usually *monoclonal* and bind to a very specific part of a protein (termed an *epitope*). Rabbit antibodies are usually *polyclonal* and less specific. Because protein sequence or structure can vary between species, antibody reactivity may vary across species (e.g., an antibody against a rat protein may *cross-react* with the equivalent mouse protein, but not the equivalent human protein).

Following primary antibody incubation, the membrane is washed to remove unbound antibody. The membrane is then incubated with a *secondary antibody*, which binds to the primary antibody and labels it. The secondary antibody is tagged with an indicator – usually the enzyme horseradish peroxidase (HRP) or a dye or fluorescent molecule. The membrane is washed again to remove unbound secondary antibody, and then subjected to a detection method, to quantify the amount of antibody binding.

(iii) Step 3, detection and quantitation

The final stage is to quantify the amount of the protein of interest, typically using chemiluminescence. The membrane is incubated in a mixture of chemicals that react with each other in a manner that is dependent on the enzyme HRP, generating a product that is luminescent (it emits light). Thus, wherever the antibody was bound to the protein of interest, that part of the membrane will light up. Quantitation of the emitted light is then performed by one of two methods:

(a) Film detection. The membrane is placed in transparent wrap (saran wrap or a clear plastic page separator) and attached to a backing plate in a metal radiography cassette. In a dark-room, a piece of x-ray film is placed into the cassette next to the membrane, typically for 1-10 minutes. Depending on the intensity of the signal

it is often necessary to try several different lengths of film exposure, to obtain a satisfactory image that is acceptable for quantitation (see discussion of dynamic range below). The film is then removed from the cassette and put through an automatic developer machine. Wherever light was emitted from the membrane, this burns a dark spot on the film. The position of the film relative to the membrane, and the position of the MW markers, are usually documented by writing directly on the developed film with a marker pen. The resulting film can be stored indefinitely, and digital images of it are usually captured with a conventional flat-bed scanner. Shown on the left is a scanned image of a piece of film, with the membrane edges marked in blue, MW markers on the left, and various other labels.

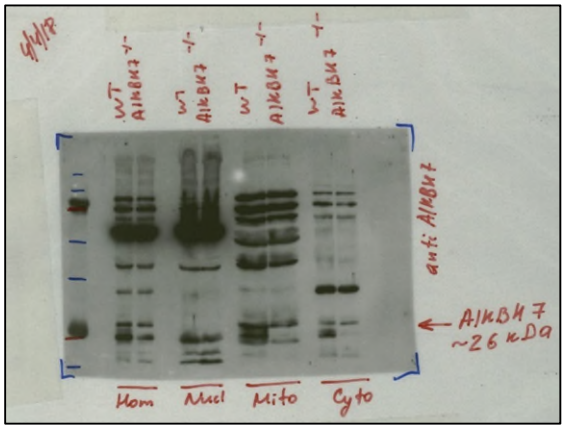

(b) Digital blot imaging. Advances in digital camera sensitivity, and the expense and toxicity of developing reagents, have led most labs to adopt digital blot imaging over the past two decades. In a digital *gel-doc* system, a series of exposures are made to collect a digital image of the chemiluminescent signal. In addition, a conventional photograph (light-field image) of the membrane is collected at the same time, so that the position of the chemiluminescent signal can be oriented to *landmarks* on the membrane (e.g., the MW ladder and the nick cut in the corner) by superimposing the images at a later time.

At this stage, the image (either physical or digital) constitutes the *original primary data* for a western blot experiment, and should be archived for future reference. The image may be altered for presentation purposes (e.g., by cropping) but the original unaltered image must be kept.

Recalling the key steps… the density of a band on a blot image is proportional to the chemiluminescent signal, which in turn reports amount of HRP-linked secondary antibody at that position on the membrane. This is determined by the amount of primary antibody bound, which is determined by the amount of the protein of interest. As such, the density of a band indirectly reports the amount of the protein of interest in the sample loaded in a given lane. Comparing band densities allows for relative comparisons between samples.

(c) Quantitation of band density

Densitometry (counting pixel density) is used to quantify the amount of a protein in each sample – i.e., to convert the blot image into numerical data. Some gel-doc systems have proprietary software, but the most commonly used software is NIH ImageJ (https://imagej.nih.gov/ij/). The software first requests to outline the lanes on the blot (see image on next page), then it plots a graph of the signal density from top to bottom for each lane. This plot can then be used to determine the area-under-the curve (AUC) for the band of interest in each lane. On the density plot, a pure white background appears as zero on the y-axis, intermediate shades (gray) in between, and a pure black band as 100.

When quantifying blots, it is rare for investigators to run authentic protein standards of known quantity alongside unknown biological samples. Thus, densitometry usually reports the relative amount of a protein between samples (e.g., sample A has three times more of the protein compared to sample B). A key issue in such comparisons is the limited dynamic range of the method (typically 10-15-fold), and a critical determinant of this range is the saturation of the blot image, determined by the exposure time. In this regard, digital imaging offers some advantages, because different exposures can be captured in composite, to extend the dynamic range of the method.

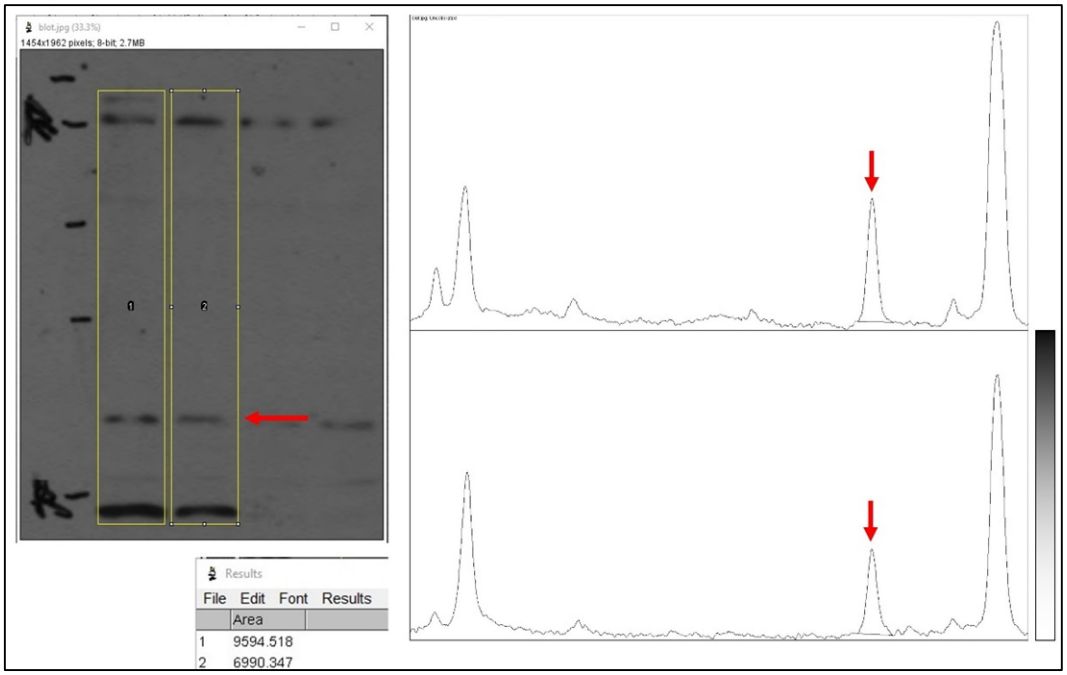

In the example at left, the blot image is shown with lanes highlighted in yellow using ImageJ software. The band of interest is indicated by a red arrow. The density plots for each lane are shown (left to right on each plot equals top to bottom on each lane). The peak corresponding to the band of interest is shown by the red arrows. Finally, the small box below the blot image shows the quantitative data from the 2 peaks (i.e., pixel density counts). In this case, the sample in lane 1 has 9594.5 units of protein, and lane 2 has 6990.3 units.

A general rule regarding saturation, is that protein bands must appear Gaussian, to qualify as non-saturated. If a band appears solid black, it will yield a density plot that is clipped or rounded at the top, indicating the signal is over-saturated. The resulting AUC will not accurately reflect the amount of protein present. In the example on the left, the peaks are rounded at the top, so the bands are not suitable for quantitation – the signal is *oversaturated* and beyond the dynamic range of the method. Many western blot images in the literature are over-saturated and unsuitable for quantitation. This issue can become especially problematic when the properties of an image are altered to render it more attractive for publication, by adjustment of the contrast/brightness, resulting in solid black bands on a solid white background.

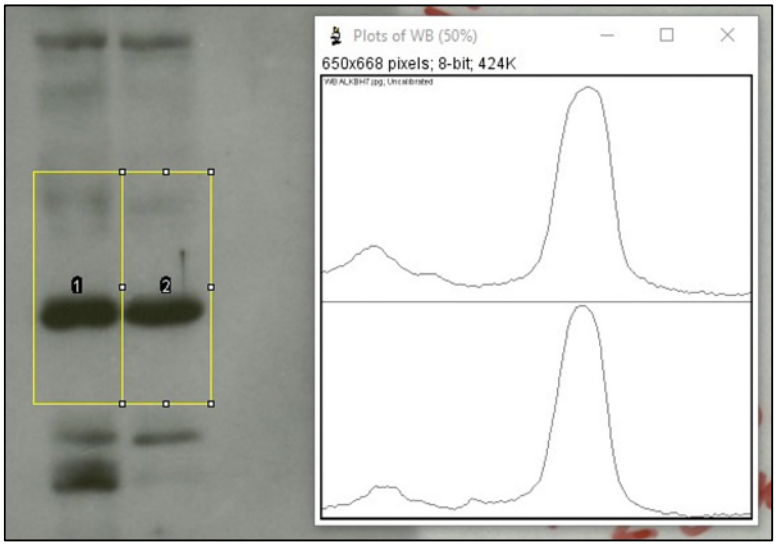

3. Loading controls

To quantify a protein of interest, it is required to normalize the signal to the amount of sample protein originally loaded on the gel. Some proteins in cells are present at relatively constant levels and are termed *housekeeping* proteins. Common examples are the enzyme glyceraldehyde-3-phosphate dehydrogenase (GAPDH) or the structural proteins β-actin or tubulin. Probing the membrane with an antibody against a housekeeping protein gives a proxy for the total amount of protein loaded in each well. An alternative normalization method is to use the Ponceau red stained membrane image (see step 2, above) to show the total protein in each lane. In best practices, a *loading control* image should be obtained from the same membrane as the one for the protein of interest. Typically, one chooses a housekeeping protein of a different MW than the protein of interest. The blot membrane is then cut into horizontal slices, and the membrane

pieces probed with different antibodies – one for the protein of interest and one for the housekeeping protein.

If the protein of interest and the housekeeping protein are too close in MW to permit cutting the membrane, it is sometimes possible to probe with one antibody, then *strip* the membrane and re-probe with another antibody. However, such procedures are difficult and do not always work, as harsh stripping procedures may damage the membrane or the protein epitopes.

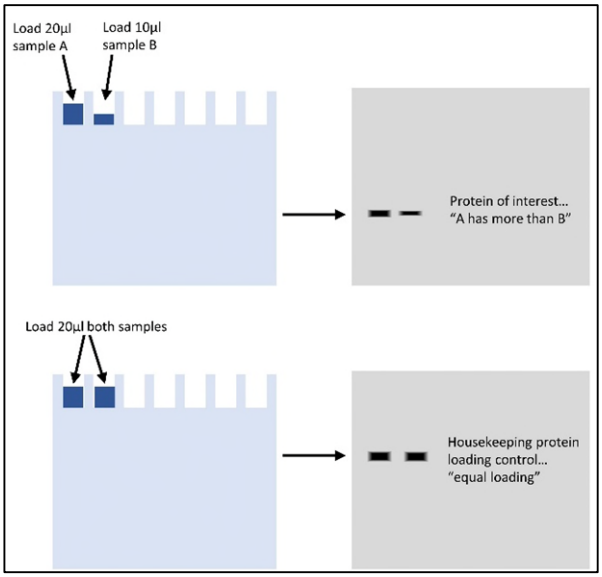

(i) Running loading controls on a separate blot

Sometimes it is simply not possible to run a loading control on the same membrane. In such cases, it can be acceptable to run a separate gel/membrane with the same samples. However, it is important to recognize that this introduces significant opportunity for misconduct. In the example shown on the right, if an investigator wanted to prove that sample A had more of a protein of interest than sample B, they could load more of sample A on the gel, and probe the blot with the relevant antibody. Then, a separate loading control blot could be constructed by loading with equal amounts of sample. *Normalizing* the protein of interest blot to the loading control blot would give the desired result, even though it is entirely fabricated.

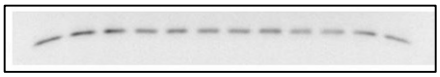

(ii) How to tell if loading controls come from the same blot?

When a blot is horizontally sliced, to probe for more than one protein, several features should be preserved across the different slices. For example, a common feature of SDS-PAGE gels is they do not run evenly, resulting in a *smile* or *frown* shape in the final bands across the gel (see example image above). If one blot image in a series has a smile/frown, then the others should have the same feature.

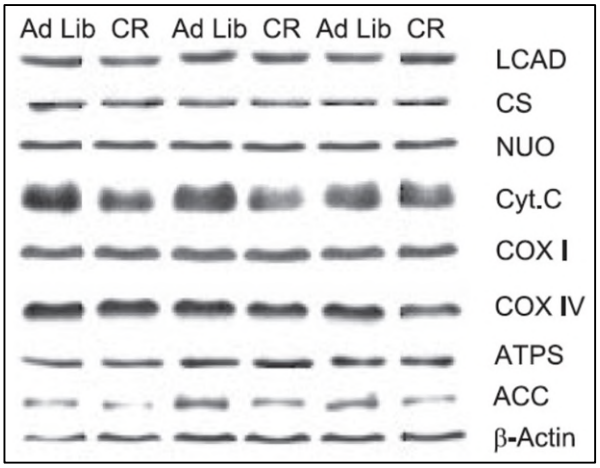

Likewise, the spacing of bands should be similar, as well as their sizing, shape and slope (left to right). In the example on the left here, we are expected to believe that the authors ran a single gel, cut the membrane up into 9 horizontal slices and probed each with a separate antibody. Clearly many of the bands show different features (e.g., the ACC bands on the right side of the image are sloped upwards from left to right, whereas the β-actin bands are cupped, and the COX-I bands have less space in between them). The more likely explanation is that these bands originated from several blots, and so the β-actin blot is not a true loading control. In such cases when it is clear that loading control images originate from separate gels/blots, it is often necessary to request full-sized original blot images as supplementary supporting data.

4. Rules and standards in preparing blot images for publication

The standard rules, accepted in the field for almost 20 years, are widely referred to as the "JCB guidelines" (Rossner & Yamada, 2004. *J Cell Biol.* vol 166, p11. https://doi.org/10.1083/jcb.200406019). The *golden rule*

is that if an image is to be manipulated, all parts must be adjusted equally. It is absolutely unacceptable to apply selective enhancement or other adjustments to only some parts of an image and not others. A critical principle from the guidelines is as follows:

> *For every adjustment to a digital image, it is important to ask, "Is the image that results from this adjustment still an accurate representation of the original data?" If the answer is "no," your actions may be construed as misconduct.*

5. Common problems & features when forensically examining blot images

(i) Letterboxing

A common practice in published blot images is *letterboxing* – cropping of blot images to show only the protein band of interest, resulting in a horizontal letterbox-shaped image. Although this has been popular for decades (and was likely bought about by space constraints in paper printed journals), many publishers now discourage it and demand original full-sized blot images in a supplementary data file.

The problem with letterboxing is it removes bands from their context. For example, best practices demand to show at least one MW marker (preferably two or more), to allow the reader to see the size of the protein. This is useful to see if the antibody is binding to something at the correct (expected) MW. Proteins are often modified so they run at a different MW, or they may be cleaved or degraded so the antibody binds to a small fragment of protein that runs further down the gel. Perhaps most importantly, MANY commercially available antibodies are known to recognize off-target proteins (e.g., Foster *et al.*, 2008, *Biochem. Biophys. Res. Commun.* vol 366, 649 https://doi.org/10.1016/j.bbrc.2007.11.154). For this reason, showing a taller swath of blot image is preferred practice and is now demanded by many journals as part of their submission process.

(ii) Blot splicing

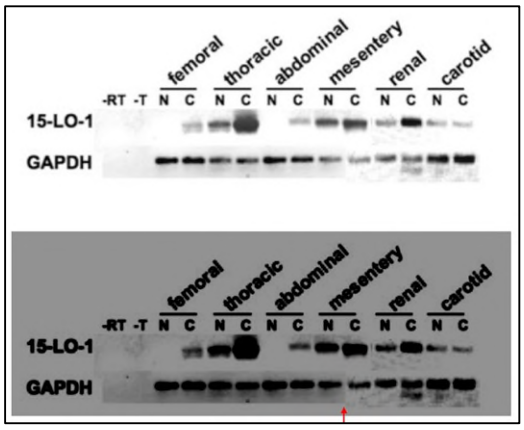

A once common practice now strongly discouraged, is the undisclosed splicing of blot images. If a gel was run with the samples in a different order than desired for publication, it may be necessary to rearrange the bands into the correct order. Doing so MUST be disclosed clearly in the figure legend, with the use of solid lines on the image itself. In the example shown at left (upper image is original published, lower image with enhanced contrast), the GAPDH loading control blot has a vertical splice seam (red arrow), while the 15-LO-1 blot does not. Ergo, the two blot images have been treated differently.

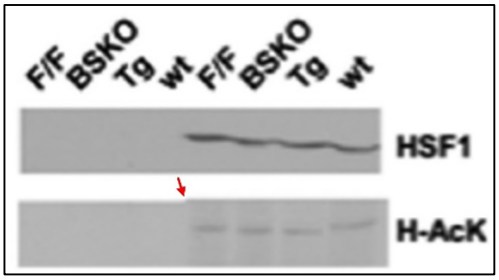

Sometimes, a giveaway sign that a blot has been spliced, is a small *notch* or *nick* along the upper or lower edge of the letterboxed image, as shown in the example on the left. Letterboxed blot images are often presented in stacks, showing several proteins from a series of conditions, with a single loading control at the bottom (see example on previous page). In such cases, it is absolutely forbidden for some blots in a stack to be spliced, and others not spliced – this would indicate the images have been treated differently.

(iii) Resolution & compression artifacts.

During or prior to the publication process, digital western blot images may go through several manipulations, with saving at each stage introducing image compression artifacts. The JPG image format is particularly *lossy*, and can introduce false edges into images that can be mistaken for splicing seams during later forensic examinations. For this reason, when forensically examining images, best practices call for obtaining the highest possible resolution image.

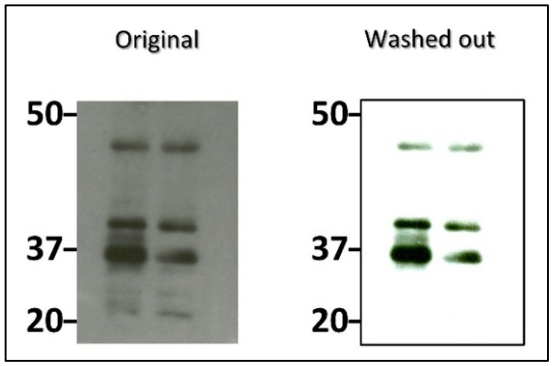

(iv) Washed out blots

In presenting blots for publication, it is common to enhance the contrast/brightness. However, when this is done to such an extent that the background of the image becomes completely *washed out*, it renders the bands disconnected from any other feature of the image (see example at right). This creates additional potential for misconduct, because it is now impossible to tell if the bands originated as they are seen, or were assembled together artificially (i.e., digitally spliced). Washed out blots are another scenario in which requesting original full-sized images is often necessary, to determine the origin of the bands.

(v) Background noise

A key determinant of whether a blot image is a genuine compound whole, or has been assembled from unrelated images, is background noise. Every blot image is unique, and carries with it pixelation and noise from the various stages of preparation. Sources of this noise include:

- Surface imperfections on glass plates used to pour the gel
- Surface imperfections and particulate matter (bits of gel) that stick to the membrane during transfer
- Fingerprints or dirt acquired during blot processing and development
- Creases in the membrane or plastic wrap overlay used during chemiluminescent development
- Surface imperfections in the x-ray film
- Dust or optical aberrations in the camera (gel-doc) imaging system or the flatbed scanner used to capture the developed film
- Residual chemicals or signal-generating features in the film cassette or imaging system (e.g., chemicals leak out of the plastic wrap).
- Dust or imperfections on rollers or other components of the automatic developer machine

Furthermore, since several stages of the blot process involve positioning things alongside each other, there are numerous opportunities for such noise and other signal imperfections to move relative to each other. For example, scanning the same piece of film on a flatbed scanner twice will give slightly different results each time, even if the film is shifted by <1mm. Likewise, when loading film into an automatic developer machine, the chances of hitting the roller in exactly the same place are essentially zero. For all of these reasons, a simple rule can be applied – NO TWO BLOT IMAGES WILL HAVE IDENTICAL BACKGROUND NOISE down to the single pixel level. Ergo, when two images do appear to share more features than would be expected by chance, it is reasonable to conclude that they share an origin in the digital realm.

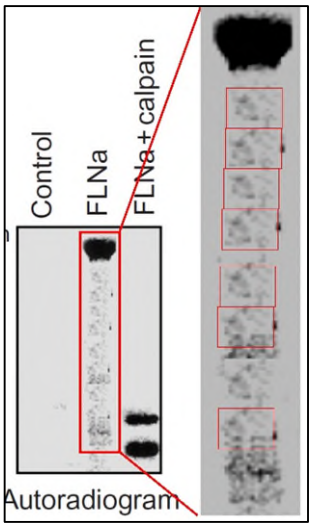

In addition to backgrounds being similar between gels or blots, it is also highly unlikely for the natural background features within a single blot image to be repetitive. In the example shown here, much of the lane beneath the prominent band at the top is filled with what appears to be random pixelated noise. However, closer examination under high contrast reveals a repeating non-natural pattern. The most likely explanation is digital *cloning* of an area of the image.

A common excuse is "there was a speck of dust on the developing machine roller", but this is untenable due to the size of the features… the cloned area is about 5 mm in size, while the rollers in developer machines are ~2" in diameter, ~6" circumference. Thus, any repeating pattern due to roller dirt should repeat every 6" on the film. This excuse also requires us to believe that the operator inserted the film at the same exact point of roller rotation, to capture that exact same speck of dust in the same position on the film.

(vi) Replication of bands and features

In the same manner that background noise should be unique to every blot image, the shape and size and characteristics of individual bands within a blot should also be unique. As such, although individual bands may appear to be similar, they should never appear identical down to the pixel level. An important feature often used to determine whether bands are identical, is the surrounding imperfections. In the example shown on the left, the bands in the top left and top right corners (red boxes) appear similar. Enlargement, enhancing contrast, and applying a color gradient map (in Adobe Photoshop software) allow the common features to be visualized more easily, leaving little doubt these are the same band.

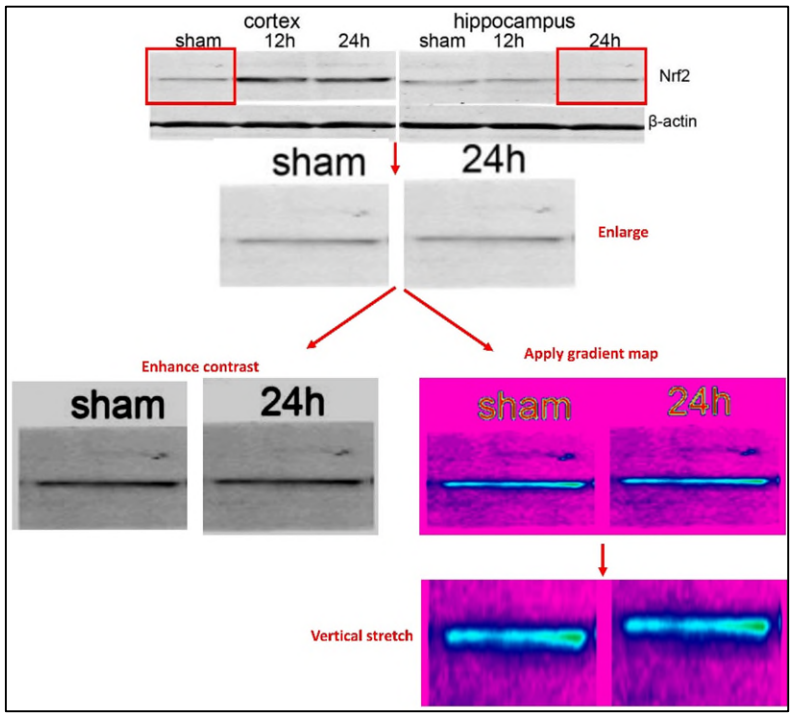

Not only are the background noise features identical, but the bands themselves have identical outlines and edges. Vertically stretching the image (bottom right) can be used to confirm that the gradient within the band itself is the same, confirming that these bands have been inappropriately duplicated. If necessary, image overlay techniques can be used to confirm if two bands are identical down to the single pixel level.

(vii) Compound images

As outlined in the section on blot development, it is common nowadays to digital imaging, in which different times of chemiluminescent exposure can be captured, alongside a light-field image of the membrane with its dye-tagged MW markers. The two images are then bought together in a composite image for presentation. Ideally, all the images including the light-field image of the membrane are captured at the same time, with

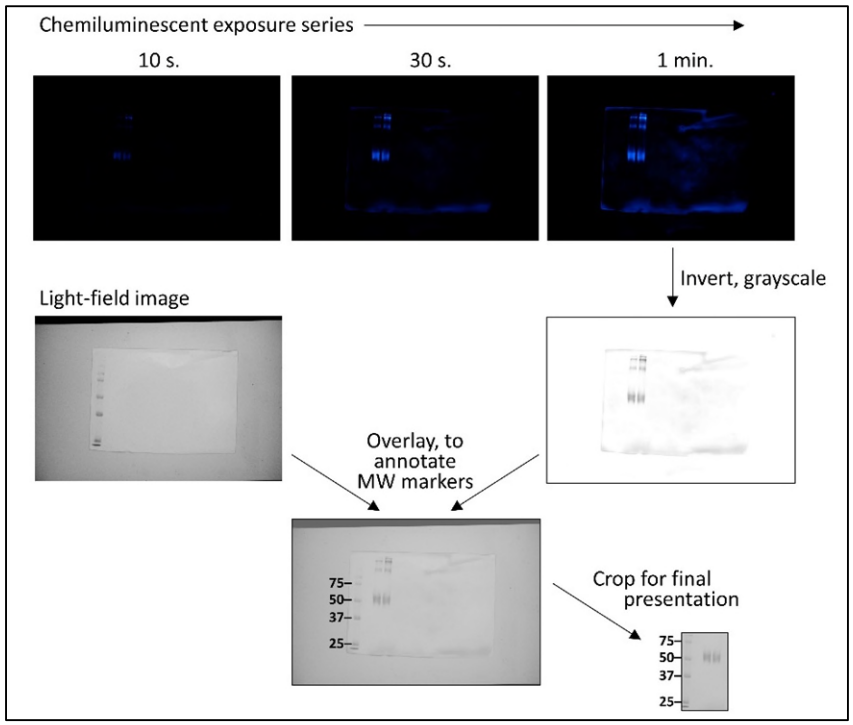

the files named appropriately and placed in a common folder. In the example shown below, a series of chemiluminescent image exposures are shown, with the light-field image of the membrane. These images are then overlaid, by making one of the images partially transparent, in order to position the MW markers. Finally, the compound image is cropped to generate the final image for presentation. It is then saved as a single layer JPG (i.e., the information about layers used to make the final image, is removed).

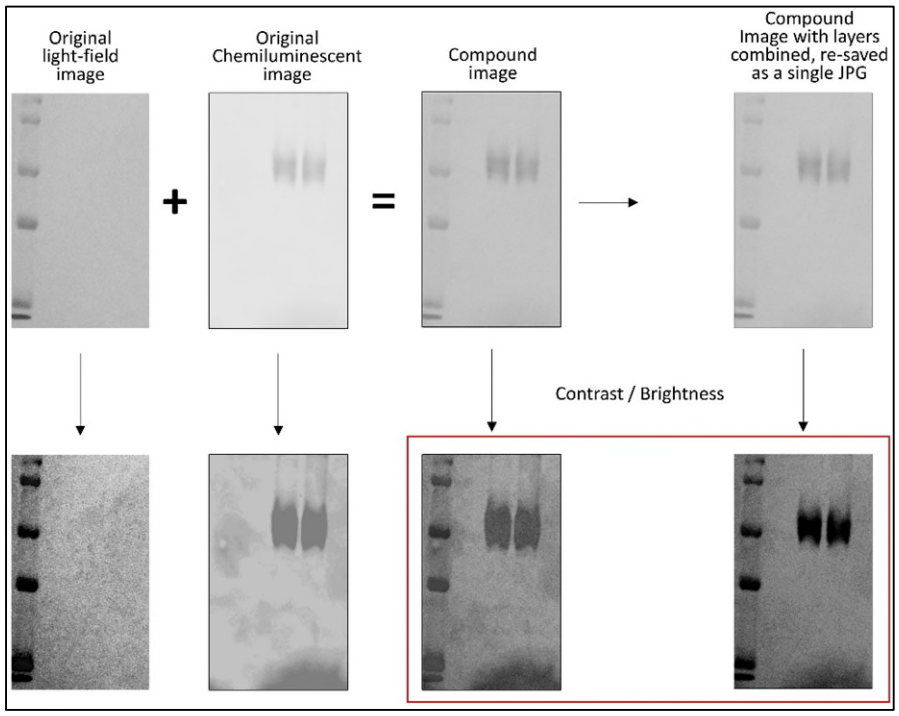

Very importantly, from a forensic perspective, ALL OF THE BACKGROUND NOISE from both layers used to make the final image, is transmitted into the final image combined. Shown on the left, in the top row are the layers generating the compound image, then the final re-saved JPG image with layer information removed.

On the bottom row are the same 4 images with brightness and contrast enhanced to show background noise. Looking at the two images highlighted in the red box, they have numerous overlapping features of background noise, but these two images are of completely different provenance! The one on the left is an overlay of two images, while the one on the right is from a re-saved JPG with no layers.

All of the background noise features from the left image (true overlay) are present in the right image (re-saved no-layers image).

The above series demonstrates that even when images are constructed from compound data collected using a digital gel-doc system, unique features from the background pixelation of each component are carried over into the final image. This series also demonstrates the important principle that background features of a final gel image are unique and should not appear across unrelated experiments. The light-field image of a membrane with MW markers (with its associated noise pattern) should not appear with the chemiluminescent bands originating from another blot membrane. Each set of bands matches to only one background, no more.

Since every image should be unique, when blot images appear to have similar background features but different bands, it is likely that one or more of the component images used to assemble the final compound image has been reused or duplicated. It is possible this can happen due to poor data management, if the component files are named in such a way that leads to confusion. However as mentioned above, best practices call for the acquisition of component images at the same time, and naming/storage in a unique folder hierarchy, to avoid such *mistakes*.

Notably, the likelihood of a series of background features appearing across multiple images is incredibly low, and becomes essentially zero when western blots are separated both by experimental type (origin and naming of samples) and time (whether the experiments were performed months/years apart). The appearance of the same background pixelation across numerous images in a data set is highly indicative of inappropriate image manipulation and would be evidence of misconduct.

2. Forensic analytical methods for western blot images

Paul S. Brookes, PhD. University of Rochester Medical Center. October 19th, 2022

Overview

Herein, I will describe standard forensic analytical methods that will be used to analyze western blot figures. Each image is subjected to the following battery of, to highlight features that can aid in determining the origins of the figure or its component parts. This document assumes prior knowledge contained in the introductory paper on western blot preparation and quantitation (the "intro-doc" #1).

1. Magnification to pixel level detail.

To aid in determining whether two images are the same or share a similar origin, magnification is often used to compare a small area of the image at the pixel level. As shown in the example below, these two images of piece(s) of film share numerous similarities (handwriting, text along the top edge, dots in the top left corner, noise in the top right edge) such that, even though one is slightly rotated relative to the other, they are images of the same piece of film

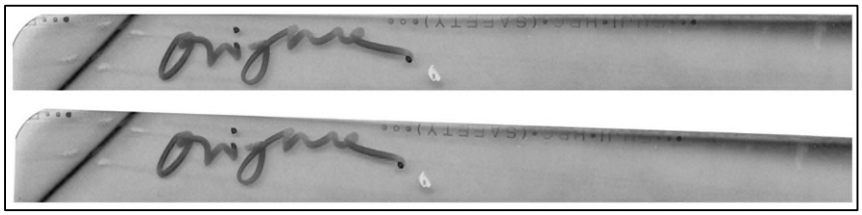

2. Brightness and contrast adjustments, PowerPoint and/or Adobe Photoshop

After importing the image into Microsoft PowerPoint (PC, Office 365, version 2207, build 15427.20210), the "Picture Format" menu is used to select the "Corrections" tab, and the contrast, brightness of the image is adjusted. Typically, this involves an increase in contrast of 50-75%, and a decrease in brightness to 30-60% vs. original. Such enhancement often reveals undisclosed seams or borders within an image, indicating where separate parts have been pasted together.

For comparisons between images, PowerPoint is the preferred tool, as it allows multiple images to be compared on a single page, and then annotations can be added to highlight common or different features.

Alternatively, for direct processing of single images (JPG, TIF, etc.), the files can be opened using Adobe Photoshop (PC, version 7.0, 2002), and the Image>Adjustments>Brightness/Contrast menu function is used to adjust the brightness/contrast of the image.

3. Curves function, Photoshop

Another method to adjust brightness/contrast within Photoshop is the "curves" function, under the Image>Adjustments>Curves menu. In a grayscale image, this function allows the user to define the output level for every input level of an image – for example one may wish to enhance dark pixels and to dull bright pixels. The unity line (straight line) leaves an image unaltered (i.e. the output image is the same as the input). The image at left shows a simple curve in red, but more complex curves may be applied, for example to enhance midtones only. Curves may also be applied in a color specific manner if the image has multiple

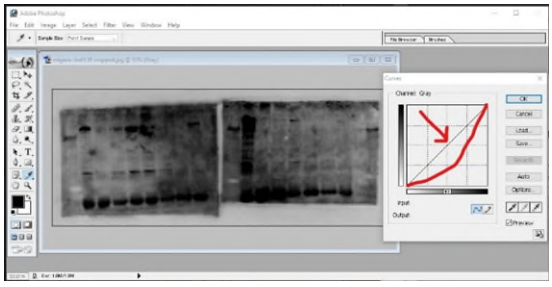

color channels, to enhance pixels of a particular hue (e.g. increase blue, decrease red).

4. Recoloring in PowerPoint or Photoshop.

In PowerPoint, in addition to brightness and contrast adjustments, a grayscale image may be recolored, i.e. application of a color mask, and this can help in visualizing similarities/differences between images, as well as highlighting edge features.

In Photoshop, a related feature is the application of a gradient map (accessed via Image>Adjustments>Gradient Map menu feature). Essentially, this feature takes each shade within an image and applies a new color to it, depending on the spectrum chosen from the menu.

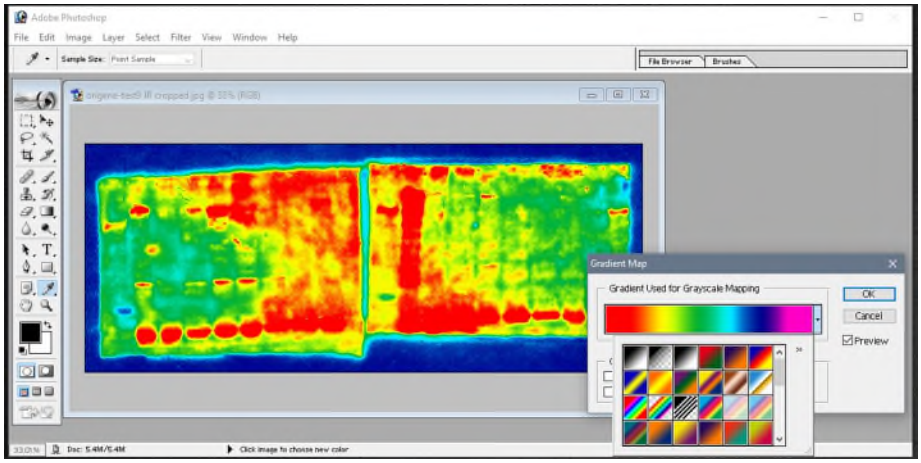

5. Histogram analysis, Photoshop

Within a western blot image, the shades available should have similar characteristics, and the range of shades within the image should be similar between similar features. That is, black should be black, white should be white. Sometimes a blot image will appear to have one or more bands where the "blackest black" is not as black as the other black bands. This can sometimes indicate that a band has originated elsewhere and has been pasted into an image with a different overall shading scheme. The histogram function is applied to different areas of an image (e.g. single bands), to show the range of shades used in that area. In the example on the left, some of the bands appear to be black, but are actually very dark gray.

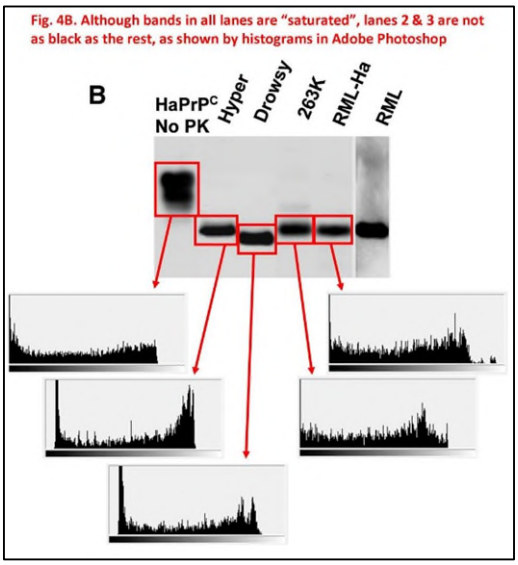

6. JPG Error Analysis, FotoForensics.

The website https://fotoforensics.com/ offers a number of image analysis tools, one of which is ELA or "Error Level Analysis". The JPG image standard includes algorithms to compress the image information, to decrease file size. It is generally understood that within a single JPG image, the entire picture should be at the same compression level. If a section of the image is at a significantly different error level, then it likely indicates this part of the image originated

elsewhere, in an image of different compression level. In the example shown below, the area appearing to originate elsewhere is highlighted by a red arrow.

7. ImageTwin

ImageTwin (https://app.imagetwin.io) is a proprietary image recognition algorithm that is driven by an artificial intelligence engine trained to spot duplications between parts of images in the life sciences. It is not yet commercially available, but I have access to a beta-testing account.

When presented with a paper that contains a large number of images, ImageTwin can look within the paper to determine if any of the components appear duplicated, and then flag them for further analysis. First, the PDF is imported, then parsed to extract images from the Figures, then the Figures are compared. The site also holds a database of all publicly available bioscience images (i.e. PubMed Central) and so can flag whether any of the images in the paper are duplicated from elsewhere in this database.

ImageTwin is more of a screening tool, and would never be used in isolation. Any images flagged by ImageTwin are always further analyzed using the other methods described here, to draw conclusions.

8. Densitometry

Using the gel densitometry feature of ImageJ software, as described in intro.doc, the density of individual bands or lanes on a blot image can be determined, to see if the numbers obtained match those presented alongside blots in graph format. This feature is also useful in determining whether a given set of blots are actually suitable for quantitation by densitometry – many published blots are over-exposed and not suitable for such analysis

Paul S. Brookes, PhD.

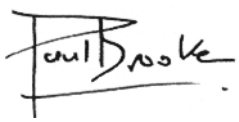

3. Analytical pipeline for western blot images

Paul S. Brookes, PhD. University of Rochester Medical Center. October 20th, 2022

Overview

This document describes a 3-stage pipeline that will be applied to the images of interest, using the analytical and forensic image analysis tools previously described (document 2).

Nomenclature - Types of Image Files

Each *analytical triad* consists of 3 related types of image files:

(A) A raw image file. The most basic images (typically JPGs) show a scan of a piece of western blot film, with no further adjustments or annotations.

(B) White box image file. In many cases, a raw image file is accompanied by a partner file with a similar file name, often with the number 11 appended (e.g., image.jpg & image-11.jpg). These images contain white boxes that appear to be superimposed upon the raw film image.

(C) Published image. The final published image is either contained within a Figure of a publication or grant application (usually a PDF file), or in a file where such figures were annotated and prepared for publication (e.g., PowerPoint file).

In certain circumstances, additional and related files may be pulled into the analysis.

Analytical Categories & Stages:

Stage 1. Reverse Analysis

In this stage, the images are analyzed in a reverse (retrograde) manner, from the final published figure (C) back to the raw image (A). The analysis asks two key questions:

Q1. Can the final published image (C) be traced to a white-box image (B) and/or a raw image (A)? The final published image will be compared with the white-box and raw images, to determine if the 3 image files have common features. These features will be used to determine if A and/or B are the source files for C. This analysis focuses mainly on the background noise and other non-data features within each image, rather than bands within the western blots.

Q2. Can the final pattern of bands in the published image (C) be traced to the source images? If it is established that a set of 3 images are related (i.e., A/B are sources for C) then I will ask whether the western blot bands in the final image can be traced to the source images. If the bands in the final image (B) can be traced to the white-box image (B), can those bands in turn be traced to raw image (A)?

Stage 2. Forward Analysis

In this stage, the analysis is performed in the forwards direction, from the raw image (A) and/or white-box image (B) toward the final image (C). The key question (**Q3**) is: **Can the source images be manipulated using accepted techniques, to recreate the final image?** The images in A/B will be adjusted using accepted techniques (i.e., applying

brightness/contrast or other adjustments to the whole image) to determine whether the final image and bands within it can be reproduced.

Stage 3. Further anomalies and notable features

In this stage, I will examine other parts of the 3 images and note any additional unusual features. This includes molecular weight markers, any apparent splicing seams, odd edges, whether final bands match densitometry data presented in the paper/grant, etc.

Documentation and conclusion: Each triad of images will yield a report, outlining answers to the 3 key questions listed above. The following colors will be used to various features between images:

RED: Feature is preserved between images

YELLOW: Feature disappears (is lost) when moving between images

BLUE: Feature appears (is gained) when moving between images

An overall conclusion statement will address whether the pattern of bands in the final published image can find provenance in the source images, and whether the scientific conclusions drawn on the basis of the published image are valid.

In drawing conclusions, the federal definition of scientific misconduct as outlined in 42 CFR § 93.103 will be used: *"Research misconduct means fabrication, falsification, or plagiarism in proposing, performing, or reviewing research, or in reporting research results."*

In particular regarding falsification, the following definition from 42 CFR § 93.103 will be used: *"Falsification is manipulating research materials, equipment, or processes, or changing or omitting data or results such that the research is not accurately represented in the research record."*

Paul S. Brookes, PhD.